The Internalist Conception of Justification

ALVIN GOLDMAN

One possible aim of epistemology is to advise cognizers on the proper choice of beliefs or other doxastic attitudes. This aim has often been part of scientific methodology: to tell scientists when they should accept a given hypothesis, or give it a certain degree of credence. This *regulative* function is naturally linked to the notion of epistemic justification. It may well be suggested that a cognizer is justified in believing something just in case the rules of proper epistemic procedure prescribe that belief. Principles that make such doxastic prescriptions might thereby "double" as principles of justification.

In the first part of this paper I contrast the regulative conception of justification with another, equally tenable, conception. Then, after noting a fundamental worry about the applicability of the regulative conception, I proceed to lay it out in more detail. The regulative justificational status of a doxastic attitude for person S at time t depends upon (a) the right set of doxastic instructions and (b) the states S is in at (or just before) t.

The regulative conception per se is neutral about the right doxastic instructions. But the question naturally arises: What *makes* this or that set of instructions the *right* instructions? The rest of the paper is devoted to this question. Two approaches are identified: *externalism* and *internalism*. Internalism takes its inspiration from a perspective that has dominated epistemology since the time of Descartes. I try to show that this perspective yields no definite or adequate answer to the question posed here; it provides no adequate conception of the rightness of doxastic instructions. This leaves externalism as the only available option, and I defend its plausibility.

Parts of this paper are positive and constructive. But the bulk of the paper is negative. It tries to undermine a classical epistemological perspective by showing that it cannot answer the question: What are the right doxastic instructions? I do not myself try to answer this question. But I do end the paper on a positive note, with a sketch of a framework within which that question may be answered.

I

In "Doing The Best One Can,"[1] Holly S. Goldman distinguishes two possible functions of a moral principle. First, it can serve as an instrument for making theoretical evaluations of actions. Second, it can serve as a device that an agent can employ to guide his/her activitives. A moral principle may not be equally useful for these two purposes. For example, the standard act-utilitarian principle—an act is right if and only if it would produce at least as much net happiness as any available alternative—is perfectly suitable (which is not to say correct) as an instrument for theoretical evaluation. It specifies conditions that determine the rightness or wrongness of an action. But this principle is not entirely suitable as an action-guiding principle. At the time of action an agent may not know which act would produce the most happiness; he may not even *believe* of any particular act that it would produce the most happiness. Thus it is unclear how to use the act-utilitarian principle to guide his conduct. In an unpublished paper,[2] Goldman suggests that we need distinct principles for decision-making purposes. One possible decision-making principle that might be associated with act-utilitarianism is to choose the act with the highest subjective probability of producing the most happiness. A different decision principle that might be paired with act-utilitarianism is to choose the act that, given your subjective probabilities, has the greatest "expected" happiness. (These two principles are not equivalent. An act with a slightly higher subjective probability than any other of producing the most happiness but a non-negligible chance of producing disaster would be enjoined by the first of these principles but not necessarily by the second.) Whichever decision principle should be associated with act-utilitarianism, the general point is that there are distinct *types* of principles: one for "theoretical evaluation" and one for practical guidance of action.

The concept of epistemic justification calls for an analogous distinction between kinds of theories or principles. On the one hand, a principle of justification might specify the features of beliefs (or other doxastic attitudes) that confer epistemic status. These features may or may not be usable by a cognizer to make a doxastic choice. On the other hand, a principle of justification might be designed specifically to guide a cognizer in regulating or choosing his doxastic attitudes. Here the criteria of justification must be ones to which a cognizer can appeal in the process of making a doxastic decision. That the theoretical and regulative functions of justification principles can be distinct emerges clearly from an account of justified belief I have proposed in another paper.[3] Refinements aside, this account—which I call "Historical Reliabilism"—says that a belief is justified just in case its causal ancestry consists of reliable belief-forming processes, i.e., processes that generally lead to truth. As a theoretical specification of epistemic status, such an account is entirely suitable. But this theory or principle cannot be used by a cognizer to make a doxastic decision; nor is it so intended. First, at the time of belief a cognizer may not know, or be in a position to find out about, the causal ancestry of his belief; and a cognizer may not know, or be able to tell, whether the processes that composed this ancestry are generally reliable. Thus there is no guarantee that a cognizer can apply the Historical Reliabilist theory to his own case. Second, Historical Reliabilism simply

is not a *rule* or *prescription* for choosing beliefs or other doxastic attitudes. It considers an *already formed* belief of a cognizer and says what features are necessary and sufficient for that belief to count as justified. It does not take a cognizer who is trying to decide which doxastic attitude to adopt vis-à-vis a given proposition and tell him what to do (doxastically speaking).

Epistemologists have been interested in theories of justification for at least two reasons. First, many have thought that a necessary condition of *knowing* a proposition is having a justified belief in that proposition. So a full analysis of knowledge requires an indication of the conditions in which belief is justified. Second, many epistemologists have been interested in "doxastic decision principles," i.e., rules for the formation of belief or other doxastic attitudes, e.g., subjective probabilities. Descartes's clearness-and-distinctness test was intended as a criterion to be used in deciding what to believe. And contemporary Bayesianism instructs cognizers to have credence functions, i.e., sets of subjective probabilities, that satisfy the axioms of the probability calculus. Many epistemologists, I believe, have conflated these two interests. They have assumed that a regulative notion of justification is the same notion of justification as the one that appears in the analysis of propositional knowledge. I think this assumption is mistaken. The best candidate for inclusion in the analysis of knowing is the Historical Reliabilist conception of justifiedness, and that notion is not a regulative one. For the purposes of the present paper, however, this issue is incidental. Here I wish to explore the idea of justification in its regulative role, whether or not it has any bearing on the concept of propositional knowledge.

I have introduced the regulative conception of epistemology by means of an analogy with ethics. It is questionable, however, whether the analogy is perfect. Ethics is largely concerned with individual actions, and actions are certainly subject to voluntary control and hence proper objects of self-guidance or regulation. But it is problematic whether doxastic states or attitudes are subject to (direct) voluntary control, and therefore problematic whether there is any point in formulating doxastic decision principles. I suspect that formation, retention, and revision of doxastic states are *not* subject to voluntary control, except perhaps in a restricted domain. Doxastic voluntarism is a dubious doctrine, however cherished it may have been by Descartes and other epistemologists. For the sake of discussion, though, let us proceed on the assumption (at least for a while) that doxastic voluntarism is true, that a cognizer can decide or choose whether to believe a given proposition at a given moment. We can then construe a principle of justification as one that instructs cognizers to adopt or retain certain beliefs (or other doxastic attitudes) in various circumstances. A principle of justification would be analogous to a moral principle that is designed to serve a regulative, or decision-making, function.

Let us be more precise about the relationship between the *justifiedness* of a belief and *doxastic decision principles* (for short, "DDPs"). We may represent a DDP as a function whose *inputs* are certain conditions of a cognizer—e.g., his beliefs, perceptual field, and ostensible memories—and whose *outputs* are prescriptions to adopt (or retain) this or that doxastic attitude—e.g., believing p, suspending judgment with respect to p, or having a particular subjective probability vis-à-vis p. Un-

less otherwise indicated, I shall here mean by "DDP" a *total* DDP, i.e., a single complete set of principles prescribing all doxastic attitudes a cognizer should have at a single time. Such a total DDP would presumably make use of a variety of different inputs, including those that pertain to perception, memory, induction, and the like.

The justificational status of believing a given proposition, say p, for cognizer S at time *t* presumably depends in part on the conditions S is in at, or just before, *t*, e.g., what evidence S possesses. But whether S is justified in believing p at *t*—whether S "ought," epistemically speaking, to believe p at *t*—also depends on the correct DDP. Assume that a unique DDP is correct, or right. Then S is justified in believing p at *t* if and only if the right DDP, when applied to the relevant conditions that characterize S at *t*, yields as output the prescription "believe p." More generally *S is justified in having doxastic attitude D vis-à-vis p at t if and only if the right DDP, when applied to the relevant input conditions that characterize S at t, yields as output the prescription* ⌜*adopt attitude D vis-à-vis p*⌝. This general relationship constitutes the basic framework of the regulative view of justification which will be presupposed in the rest of our discussion.[4]

Given this approach to justificational status, several questions obviously become paramount. First, what are admissible sorts of input conditions for a DDP? Which states of a cognizer are relevant to the justificational status of a doxastic attitude? Second, are we right in assuming that there is a uniquely correct DDP? Third, if this assumption is correct, what *makes* a certain DDP right, or correct?

II

Let us begin with the question concerning admissible inputs. In illustrating input conditions, I mentioned various cognitive states of a person, e.g., his beliefs and ostensible memories. It is worth asking, however, why the relevant input conditions should be cognitive states, or, for that matter, why they should be *states* of the person at all. Why could input conditions not be states of the world, or the external environment? On purely formal grounds, the following seems to qualify as a DDP: "For any proposition p, if p is true, then believe p (at any time *t*)." The input conditions for this DDP are not states of the cognizer; rather, they are the truth-values of the various propositions, or the "states of the world" that make these propositions true or false. But why not allow a DDP with input conditions of this kind? Admittedly, this is an intuitively inappropriate DDP. But what exactly makes it inappropriate?

The answer is straightforward. If a DDP is to be actually *usable* for making deliberate decisions, the conditions that serve as inputs must be *accessible* or *available* to the decision-maker at the time of decision. The agent must be *able to tell*, with respect to any possible input condition, whether that condition holds at the time in question. Now if the truth-value of any random proposition is a possible input condition, a cognizer would have to be able to tell, with respect to any such proposition, whether or not it is true. This requirement is not satisfied. Hence, the general class of truths and falsehoods cannot serve as the appropriate domain (i.e., input conditions) for a DDP.

We now see why a person's current cognitive states are a plausible class of input conditions. It is plausible to hold that for such states a person *can tell*, at any moment, exactly which of them he is in at that moment. So these input conditions would satisfy the requirement of being "accessible" or "available" to the decision-maker. But what exactly do we mean in saying that a person "*can tell*" with respect to a given condition whether or not that condition obtains? Here is a reasonable answer: "For any person S and time t, if S asks himself at t whether condition C obtains at the time in question, then S will believe that condition C obtains then if and only if it does obtain then."[5]

Notice that *past* cognitive states do not satisfy this constraint. It is not true, in general, that if I ask myself at time t whether or not I was in a certain cognitive state G at an earlier time t_0, then I will believe that I was in G at t_0 if and only if I was. I may forget or misremember my past cognitive states. For this reason, cognitive *ancestry* is not among the conditions that can serve as a DDP input. Thus a "historical" theory is excluded as a *regulative* theory of justification.

It is worth exploring some other consequences of our constraint on input conditions. Epistemologists commonly include logical relationships in their epistemic rules. For example, a rule might say: "If you are justified in believing Q, and Q *logically implies* P, then believe P." One input for this rule is the obtaining of a logical implication. But according to our constraint this is not an admissible input. It is not in general true that a person *can tell*, for any propositions Q and P, whether or not Q logically implies P. It appears, then, that many favorite examples of epistemic rules may not be legitimate (portions of) DDPs.

The rule in the previous paragraph poses another question about admissible input conditions. Is the justificational status of a doxastic attitude a legitimate input condition? Is *being-justified-in-believing Q*, as opposed to merely *believing Q*, an admissible input? If we admit (regulative) justificational status as an input condition, we have a threat of circularity in our theory. The aim of specifying a class of inputs and a correct DDP is to provide a theory of justification. If (the notion of) justificational status itself appears in the input conditions, our account would seem to be circular.

The charge of circularity should not be leveled too hastily. If a recursive account of justification were given, it would be unobjectionable to have a *recursive clause* by which the justificational status of believing Q could help determine the justificational status of believing P.[6] However, such a recursive account would also need *base clauses*, and to avoid circularity these base clauses would have to specify non-justificational conditions—substantive or "factual" conditions—for justificational status. So the justificational status of believing P would ultimately be traceable to these substantive conditions, which may be viewed as the relevant inputs. Thus it is appropriate to add the following restriction to our constraint on inputs: inputs must be purely factual, non-epistemic conditions.

This is all I shall say about admissible inputs for a DDP. Let us now turn to the question of what makes something the right DDP, and whether, indeed, a uniquely correct DDP can be demanded.

III

The choice of a DDP clearly depends on the *goals* of cognition, or doxastic-attitude-formation. A very plausible set of goals are the oft-cited aims of *believing the truth*—as much truth as possible—and *avoiding error*. (Some alternative goals will be examined later.) These twin desiderata, however, tend to compete with one another. A "conservative" DDP prescribes more suspension of judgment than does a "venturesome" DDP. Greater conservatism would tend to produce less false belief, which is good, but also less true belief, which is bad. Which of two such DDPs is preferable, on the whole, is a function of how the totality of true belief is weighted as compared with the amount of error. One view would be that a single false belief outweighs a tremendous amount of truth. Another view would be that there is as much positive value in a single (modest) truth as there is negative value in a single (modest) error. For my purposes, this knotty issue can be sidestepped. The issues I wish to raise in this paper are independent of the weighting problem. So let us proceed on the assumption that *some* combination of true belief and error avoidance is what we seek in a DDP.[7]

Given the aim of true belief and error avoidance, the right DDP is apparently one that would produce optimal results in terms of true belief and error avoidance. It is the DDP that would have such optimal results in the long run for the sum-total of cognizers. Or, assuming that what is best for one (human) cognizer is best for others, the right DDP is the one that would produce optimal results for any cognizer taken singly. It is the DDP that God in his omniscience would recommend.

Unfortunately, the foregoing characterization of the right DDP ignores a crucial aspect of traditional epistemology. The foregoing conception rests on an "*externalist*" perspective: the perspective of a Godlike observer who, knowing all truths and falsehoods, can select the DDP that optimally conduces to true belief and error avoidance. Traditional epistemology has not adopted this externalist perspective. It has been predominantly *internalist*, of egocentric. On the latter perspective, epistemology's job is to construct a doxastic principle or procedure *from the inside*, from our own individual vantage point. To adopt a Kantian idiom, a DDP must not be "heteronomous," or dictated "from without." It must be "autonomous," a law we can give to ourselves and which we have *grounds* for giving to ourselves. The *objective* optimality of a DDP, on this view, does not make it right. A DDP counts as right only if it is "certifiable" *from within*.

To illustrate the point, suppose a DDP were proposed that consisted in a very, very long list of propositions to be believed: propositions about individual events and states of affairs, laws of nature, and so on. Belief in these propositions is prescribed *unconditionally*, independent of the cognitive states of the agent. In short, the input conditions for this DDP are the null set. Further suppose that all the prescribed propositions in this long list are *true*. Does that make this DDP a strong candidate for the right DDP? Not at all, according to internalism. This is a DDP that a Godlike observer might give us, not the sort we can legitimately give to ourselves. More cautiously, if we *are* in a position to give that DDP to ourselves, it must be because we have used some other, more fundamental, DDP to ascertain the relevant

set of truths. It is that more fundamental DDP which ought to be proposed as the genuinely correct DDP.

The foregoing example, it must be conceded, has oddities that are irrelevant to the point at issue. The DDP in question is counterintuitive because it contains no *general* instructions, no precepts for generating new beliefs from old, no provisions for learning from experience. All these features would naturally be expected in a good DDP. For this reason, let us illustrate the internalist idea with another, more familiar, example: the problem of induction.

Assume for the sake of argument that there is a unique inductive rule that would actually be optimal for predictive puposes. (Specification of an inductive rule will probably require, to meet Goodmanian concerns, either an enumeration of projectible predicates or a set of rules of projection. Assume these are built into the rule in question.) The mere fact of optimality, however, would not be regarded by internalism as a solution to the problem of induction. According to internalism, an inductive rule is right only if we can *justify* the claim that it is optimal, i.e., only if we can *show* that it will lead to truth, or will probably lead to truth, or that it (alone) meets some weaker (e.g., Reichenbachian) test of optimality. *De facto* optimality does not satisfy the internalist. A rule that just *happens* to be optimal is not right unless we *know*, or are *justified* in believing, that it is optimal.

Thus far I have given a rather vague characterization of the internalist conception of justifiedness. In the remainder of the paper I want to explore this conception critically. I shall argue that the internalist conception is either fundamentally confused or unfulfillable. Either (A) there is no definite and acceptable set of conditions that articulate the vague idea of the internalist, or (B) although such a definite set of conditions can be specified, there is nothing—no DDP, or rule of justification—that meets these conditions. In short, what appears to be a comprehensible and attractive conception of justification melts away when examined carefully. Finally, I shall argue that externalism provides a perfectly satisfactory conception of justifiedness.

IV

The crucial question for internalism is: What is the right DDP? We shall not try to pinpoint the particular DDP that is right according to internalism (nor the one that is right according to externalism). Rather, we shall try to see whether there is a definite and acceptable set of conditions that determine what internalism would *count* as the right DDP. We can indicate the sort of condition being sought by formulating the condition appropriate to *externalism*, viz. (1):

(1) DDP X is right if and only if: X is *actually* optimal.

By 'optimal' I mean, of course, optimal in producing true belief and error avoidance. As indicated earlier, the exact weighting of true belief and error avoidance is here neglected.

For reasons given above, (1) does not satisfy the internalist. What, then, *would* do so? Judging by our earlier characterization of internalism, the condition that seems to capture the internalist's conception is this:

(2) DDP X is right if and only if: we are *justified* in believing that X is optimal.

There is a fatal problem with (2), however. It uses the notion of justification! As indicated previously, the aim of a theory of justification is to assign justificational status in *non-justificational* terms. In particular, the aim of a regulative theory of justification is to provide instructions concerning doxastic attitudes that do not presuppose the prior existence, or establishment, of any such prescriptions. If the (regulative) notion of justifiedness is allowed in the condition(s) for the rightness of a DDP, this requirement will clearly be violated. For the rightness of a DDP is one of the two basic components of the theory of justifiedness. In short, proposal (2) is blatantly circular, and inadmissible on purely "formal" grounds.

It will do no better to say, in place of (2), that a DDP must be "certifiable" as optimal, or that we must have "grounds" for believing it to be optimal. Terms like 'certifiable' and 'grounds' are themselves epistemic terms, roughly synonymous with 'justified'. All such terms are equally inadmissible for the purposes at hand. Unfortunately, the initial sketch of the internalist conception utilized these very terms. This suggests that (2) correctly expresses the intuitive idea behind internalism, unacceptable as that idea turns out to be. But there may be other, unobjectionable ways of fleshing out internalism. Let us explore some further possibilities.

In place of (2) we might try (3):

(3) DDP X is right if and only if: we *believe* that X is optimal.

Clearly, (3) avoids the formal problem facing (2). But is it at all plausible? I think not. Suppose we believe some particular DDP X to be optimal for (intuitively speaking) very bad reasons, or for no reasons at all. We may believe it to be optimal out of wishful thinking, sheer confusion, or mere hunch and guesswork. We may believe it to be optimal simply because it pops into our heads, or comes to us in a dream. Is internalism committed to saying, in such circumstances, that X is really right? That doxastic attitudes *ought* (epistemically speaking) to be formed in accordance with X? Surely not. Internalism presumably does not want the rightness of a DDP to be determined by sheer eccentricity, frivolous reasons, or happenstance.

These considerations are just the sorts of considerations that motivate proposal (2). But we cannot go back to (2). Is there anything similar to (2) that does not violate the formal restriction on which (2) itself founders? It may seem promising to recall that a non-regulative theory of justified belief is available: Historical Reliabilism. It may be suggested that no circularity would be involved if we used the non-regulative notion of justifiedness in our theory of the regulative notion. More important, since Historical Reliabilism is a theory formulated in non-epistemic terms, why not use the substance of this theory without employing the term 'justification' (or any of its cognates)? This would yield the following:

(4) DDP X is right if and only if: (A) we believe that X is optimal, and
(B) this belief was caused by reliable cognitive processes.

Unfortunately, (4) founders on another restriction, a restriction peculiar to inter-

nalism. The basic idea of internalism is that there should be guaranteed epistemic access to the correctness of a DDP.[8] No condition of DDP-rightness is acceptable unless we have epistemic access to the DDP that in fact satisfies the condition, i.e., unless we can tell which DDP satisfies it. The internalist's objection to externalism's condition of rightness, i.e., actual optimality, is precisely that cognizers may have no way of telling which DDP satisfies it. Internalism's *own* condition of rightness must, therefore, be such that any cognizer *can tell* which DDP satisfies it. But this restriction is not met by (4). In general, we are not in a position to tell how some belief of ours was caused; nor is it guaranteed that we can tell which of our cognitive processes are reliable and which are not.

In addition to the foregoing objections (which are conclusive enough), there are other problems with (2), (3), and (4). A problem common to all three is the use of the term 'we'. To whom does this pronoun purport to refer? To *everyone*? This implies, in the case of (3) and (4), that a DDP is right (according to internalism) only if *everyone* believes it to be optimal. But surely such universal consensus is hard to come by and unreasonable to require.

How is this difficulty to be met in any future proposal? Should 'we' be taken to refer to a *majority* of people? To a *plurality* in favor of a single DDP as compared with any other DDP? Neither of these suggestions is attractive. A more promising solution is to *relativize* DDP-rightness to a cognizer (and a time). This would yield the following analogues of (3) and (4):

(3*) DDP X is right for S at t if and only if: S believes at t that X is optimal.

(4*) DDP X is right for S at t if and only if: (A) S believes at t that X is optimal, and
(B) this belief was caused by reliable cognitive processes.

The relativization solution is a serious step that should not be taken lightly. But let us defer an examination of that issue for a moment.[9] There is still another problem that would face (3*) and (4*) as they stand, a problem related to the one that relativization was intended to solve. Relativization was intended to meet the objection that the same DDP may not be believed *by everyone* to be optimal. But we must also note that some people may not believe of *any* DDP that it is optimal. Indeed, since the notion of a DDP is an abstruse notion, unlikely to be familiar to anyone but a philosopher, it is highly probable that *most* people will not believe of any DDP that it is optimal. It will follow from (3*) and (4*) that no DDP at all will be right for the vast majority of people. This would seem to imply that no doxastic instructions govern these individuals, so that there are no doxastic attitudes that they are justified (or unjustified) in choosing. Certainly this cannot be intended by internalism. Thus what internalism needs is a condition of rightness that does not require *actual belief* by a cognizer in the optimality of some DDP.

To meet this requirement the internalist might say that DDP X is right for person S at time t just in case S *would* believe X to be optimal, given S's antecedent mental states. But would . . . *if what*? If S formed this belief in accordance with

the right DDP? Obviously the internalist cannot say that, for it would be a flagrant circularity. What might be said, however, is that S would believe X to be optimal if S used X itself. In other words, X is right for S at *t* just in case X is "*self-prescribing*" for S at *t*. This may be written as (5):

(5) DDP X is right for S at *t* if and only if: if X is applied to S's (relevant) inputs at *t*, X prescribes belief in ⌜X is optimal⌝.

One problem with (5) is that it may fail to meet the internalist's restriction that a cognizer should be able to tell which DDP meets the condition. If a DDP is very complex, it may be difficult to tell what it prescribes, in particular, whether or not it prescribes belief in its own optimality. A second problem is even more critical. Two or more incompatible DDPs can each satisfy (5) for the same person and time. Let Y be a DDP that, among other things, prescribes use of the "scientific method" (assuming there is a unique such method). Let Z be a DDP that, among other things, prescribes belief in accordance with ostensible revelations (rather than the scientific method). Then it may well happen that, given S's inputs at *t*, Y prescribes belief in ⌜Y is optimal⌝ and Z prescribes belief in ⌜Z is optimal⌝. According to (5), *both* Y and Z are right for S at *t*. But since they are incompatible, they will generate some incompatible doxastic prescriptions. Which of these competing prescriptions should be followed? Which should be used to assess the justificational status of S's doxastic attitudes? (5) cannot answer these questions and is therefore unacceptable.

V

Let us return now to the issue of relativization. It follows, of course, from the strategy of relativization that a DDP which is right for one cognizer may not be right for another. Is such "*epistemological relativism*" acceptable to internalism? It is clear, I think, that internalists have not generally intended any such relativism. They have usually assumed that some uniform set of doxastic principles should govern all human cognizers. Foundationalists think that foundationalism (insofar as it generates doxastic prescriptions) is right for everyone. Coherentists think that coherentism is right for everyone. Bayesians think that Bayesianism is right for everyone. So we will need some special and strong argumentation to sustain the idea that internalism should incorporate relativism.

One possible line of argumentation would attempt to draw an analogy with the distinction in ethics between *objective* and *subjective* rightness.[10] An act is said to be objectively right just in case it actually satisfies the conditions for moral rightness. An act is said to be subjectively right just in case, roughly speaking, the agent's beliefs or evidence concerning the circumstances or consequences of the act suggest that it is objectively right. Subjective rightness involves relativization to the beliefs or evidence of the agent at the time of action. Now it might be suggested that the moral status of objective rightness is analogous to the *ex*ternalist conception of DDP rightness and that the *in*ternalist conception of DDP rightness is analogous to the moral status of subjective rightness. Since the moral notion involves relativization, it is proper for the epistemic notion to involve relativization as well.

Apart from the issue of relativization, the idea of subjective rightness in ethics may hold promise of assistance in our present predicament. After all, the notion of subjective rightness in ethics is usually thought to be tolerably clear. Why not borrow its analysis for the purpose of stating conditions of DDP rightness?

Regrettably, matters are not so simple. There seem to be two basic strategies for analyzing subjective rightness in ethics: (A) Action A is subjectively right if and only if the agent *believes* of A that it is objectively right; or (B) Action A is subjectively right if and only if the agent is *justified* in believing of A that it is objectively right. We may call these the "*doxastic*" and "*epistemic*" approaches, respectively. Neither of these approaches works well when applied to our present topic. When transferred to our topic, the doxastic approach yields (3) or (3*), which proved quite unsatisfactory. (Some of the same criticisms might apply to the *moral* notion as well.) Furthermore, the epistemic approach would yield our proposal (2) (or a relativized version of (2)), which is wholly inadmissible. The moral theorist is entitled to use an epistemic notion to explicate *his* concept of subjective rightness, but *we* are not entitled to use an epistemic notion, for our *explicandum* is itself epistemic.

The ethicist's notion of subjective rightness, then, proves to have less utility for our purposes than we might hope. Still, it may help support the idea of relativization. What we need, however, is not just the analogy from ethics but a particular example that motivates this idea in the epistemic sphere. Here is an example designed to serve this purpose.

Consider two men who belong to different cultures in different historical periods. The first belongs to a prescientific or early scientific community, in which precise methods of experimentation and statistical techniques have never been dreamed of. The second belongs to a scientifically advanced culture and has personally been trained in methodological and statistical techniques. Imagine that each of these men happens to entertain the same scientific hypothesis, H, and that each has the same observational evidence that bears on H. According to the DDP believed by the first man to be optimal, H should be believed. According to the DDP believed by the second man to be optimal, H should not be believed. Now suppose each man adopts the doxastic attitude vis-à-vis H prescribed by his favored DDP. Is each not *justified* in adopting this doxastic attitude? In particular, the member of the prescientific culture isn't justified—at least "subjectively justified"—in adopting his doxastic attitude? After all, the sophisticated techniques of advanced science are, by hypothesis not yet invented. He can hardly be faulted for proceeding as best he can, by his own lights. Is not the DDP he uses at least "right for him?"

Although this example is plausible, it is not wholly convincing. A defender of epistemological objectivism might reply as follows. "It is admittedly *excusable* or *pardonable* for the member of the prescientific community to accept his own DDP—call it W—and for him to form the doxastic attitude prescribed by W. But to say that this is *excusable* is not to say that the resultant doxastic attitude is *justified*. A doxastic attitude is justified only if it is enjoined by proper procedure, and the description of the case presupposes that DDP W is not proper procedure, i.e., is not the right DDP. To be sure, it may be difficult for this cognizer, in his historico-cultural setting,

to identify the right DDP. This is why we would not *blame* him for his doxastic act. Still, there is no need to consider him *justified* in believing H. Nor is there need to say that W is in *some* sense a 'right' DDP. The only right DDP is the objectively right DDP. Even internalism should insist on an objectively right DDP, whether or not it is equivalent to the 'optimal' DDP."[11]

I am inclined to endorse this response of the objectivist and conclude that the example in question should not impel the internalist to be a relativist. On the other hand, even if the internalist wants to be a relativist, we have not yet identified conditions that express an adequate version of internalism, either of a relativistic or a non-relativistic variety. Thus we have yet to make cogent sense of the vague intuitions that motivate internalism. Thus far internalism is a mere will-o'-the wisp.

VI

At this juncture the internalist may object that I have ignored the most obvious and straightforward way of specifying his conception. At least since the time of Descartes there has been a reasonably well-defined idea of the "*internal standpoint*" in epistemology. This is the idea of a certain epistemological *starting point*, a position from which all doxastic decisons are to be made. To make doxastic decisions, however, one needs a DDP, and so it is natural to say that the DDP must also be chosen from the same "internal" starting point. This was Descartes's own practice, for his criterion of clearness-and-distinctness was proposed and argued for from this very vantage point.[12] We have, then, a natural way of formulating the internalist's condition for the rightness of a DDP:

(6) DDP X is right if and only if: X is the proper DDP to choose if one chooses a DDP from the internal standpoint.

So formulated, internalism certainly contrasts with externalism, for there is no guarantee that the DDP properly chosen from the internal standpoint would be identical with the optimal DDP.

Is (6) a clear and definite condition for DDP-rightness? That depends on how clear and definite we can make the idea of the internal standpoint and the criteria of "proper choice." I shall argue that, however clear and definite we can make this idea and these criteria, there is no unique DDP that can be generated.[13] In short, even if (6) is a condition that satisfactorily expresses the spirit of internalism, it does not select or determine a preferred DDP. In this sense the internalist conception of a right DDP is "*unfulfillable*."

A merely apparent difficulty for choosing a DDP from the internal standpoint is the threat of a vicious regress. The choice of a DDP requires a criterion of choice, and it might be supposed that this criterion must itself be a DDP: a meta-DDP. But how is *that* DDP to be chosen? By means of a meta-meta-DDP? And so on ad infinitum? This worry rests on a confusion. A DDP is not a proposition but a *policy* or set of conditional *prescriptions*. (*That such-and-such a DDP is optimal* is a proposition; but the DDP is not a proposition.) Since a DDP is not a proposition, the

adoption of a DDP is not the adoption of a *doxastic* attitude. Since no *doxastic* attitude is being chosen, no DDP is required. To be sure, some criterion of choice is needed, which does, as we shall see, pose problems. But the criterion in question may be a familiar decision criterion, such as maximizing expected value, not a DDP. Thus no infinite regress is launched.

A somewhat analogous problem, though, is the central problem that confronts the attempt to choose a DDP from the internal standpoint. Although the choice of a unique DDP does not require a prior (meta-) DDP, it *does* require antecedent doxastic attitudes on which to base a choice. To assess the probable consequences of this or that policy of forming physical-world beliefs, for example, one needs some doxastic attitudes toward propositions that describe typical events in the physical world and relationships between the physical world and one's mental states. But such doxastic attitudes are absent in the internal standpoint.

To clarify this point, let us be more precise about the nature of the internal standpoint. As indicated, this standpoint is supposed to be an epistemological starting point. At such a point one should be epistemologically neutral or uncommitted. More specifically, *doxastic* neutrality is required. The whole point of choosing a DDP, after all, is to license selected doxastic attitudes. Until a DDP has been properly chosen, no doxastic attitudes can be licensed, at least no doxastic attitudes toward epistemologically problematic propositions. So the internal standpoint must disallow the use of, or appeal to, prior doxastic attitudes we may happen to have. As in Rawls's "original position," the internal standpoint is a perspective in which principles must be chosen from behind a "veil of ignorance."

I said that no appeal may be made to doxastic attitudes toward epistemologically "problematic" propositions. Which propositions are problematic? Virtually all contingent propositions have generally been regarded as problematic, and doxastic attitudes toward these should certainly be barred from the internal standpoint. A possible exception is first-person-current-mental-state propositions, which are often regarded as unproblematic although they are contingent. We may allow beliefs in these propositions to be used in the internal standpoint, just for the sake of argument.[14] It is unlikely that this will substantially assist the internalist's project anyway.

The foregoing characterization of the internal standpoint is obviously modeled on Descartes's procedure. Yet Descartes himself is not sufficiently thoroughgoing. He advocates the policy of *suspending judgment* on problematic propositions at the epistemological starting point. But suspension of judgment is itself a kind of doxastic attitude. What justifies Descartes in adopting *that* attitude toward the propositions in question? He operates on the assumption that when evidence is inconclusive you should suspend judgment. But this in itself is a partial DDP, and, like any DDP, should first be argued for. Before the selection of a DDP we cannot say that the proper attitude to adopt is suspension of judgment. This is not doxastic neutrality.

It must be admitted, though, that suspension of judgment may be the best approximation to doxastic neutrality, or doxastic nullity. If you have a proposition in mind at all, how else can you avoid commitment? So let us permit the use of suspension of judgment in the internal standpoint. However, suspension of judgment must

not be equated with a subjective probability of .5. *Such* a doxastic attitude certainly would be prejudicial.

Let us now see how the absence of usable doxastic attitudes in the internal standpoint makes it impossible to choose a unique DDP. The most popular criterion of choice in decision theory is *maximizing expected value*. Suppose that this criterion is used in trying to make the choice. In the present case maximizing expected value means selecting the DDP that has the greatest expected value as measured by true belief and error avoidance. For each DDP the "expected" verific consequences of adopting it would be computed and then the DDP with the greatest expected value would be chosen. The rub, however, is that the notion of "expected" value is partly a function of *subjective probabilities*. The "expected" value of adopting a given DDP is a function of one's estimate of that DDP's leading to truth and away from error. But we have already seen that no such subjective probabilities vis-à-vis contingent subject matter—at least physicalistic subject matter—can be appealed to in the internal standpoint. Hence the criterion of maximizing expected value cannot be used in the internal standpoint to select a unique DDP.

A concrete example may help. Suppose you are trying to decide whether to choose a DDP that incorporates the rule: "Whenever you experience a certain sequence of sensations, believe the proposition: 'There is a doorknob before me.'" To make this decision you need an assessment of the consequences of including this rule. Will adoption of it tend to produce true beliefs or errors? What proportion of each? Unless you have a (doxastic) estimate of these outcomes, you cannot assess the expected value of this rule. But any such estimate is a doxastic attitude toward a *contingent* proposition and is therefore barred from the internal standpoint. This is so even if phenomenalism is true. Even if phenomenalism is true, any given sequence of sensations is compatible not only with there being a doorknob before you but also with, for example, your hallucinating the feeling of touching a door knob. Since the sequence of sensations is compatible with both states of affairs, an estimate of the verific consequences of adopting the rule in question will depend upon an estimate of the frequency with which you will really encounter doorknobs as opposed to merely hallucinate such encounters. But no such estimates are permitted in the internal standpoint.

My argument for the impossibility of choosing a DDP from the internal standpoint has thus far rested on a particular decision criterion: maximizing expected value. But decision theory has a larger inventory of decision rules, and some of these might dictate a choice without the indicated doxastic attitudes. Consider the criterion of maximin, for example. This criterion tells a decision-maker to choose an option (here, a DDP) whose worst possible outcome is at least as good as the worst possible outcome of any other option (DDP). Trivially, the worst outcome for a DDP that enjoins any beliefs at all is the outcome where all enjoined beliefs turn out false. So the worst outcome of a DDP that never enjoins belief at all—the skeptic's, or agnostic's, DDP—is clearly the best of the worst outcomes. Thus it might appear that maximin *would* generate the choice of a unique DDP from the internal standpoint, namely, the skeptic's DDP.

But why use the maximin criterion when choosing a DDP from the internal standpoint? Why not select, say, the maximax criterion instead? Maximax tells you to choose an option (here, a DDP) whose best possible outcome is at least as good as the best possible outcome of any other option (DDP). Obviously, maximax would generate a very different choice of DDP than maximin would, if it generates any unique choice at all.[15] Maximax would mandate a DDP that prescribes belief most liberally (indeed, profligately), since the best possible outcome of such a DDP—viz., all the enjoined beliefs being true—will be unsurpassed.

Perhaps *some* situations could dictate a preference for maximin over maximax and other situations the reverse preference. But the internal standpoint permits no preference whatever. *Ex hypothesi*, the internal standpoint contains no doxastic attitudes toward (problematic) contingent propositions, so there is no basis for either optimism or pessimism in one's choices. Without a basis for selecting a unique criterion of choice, no unique DDP can be generated from the internal standpoint.

Have I been too restrictive in excluding doxastic attitudes toward virtually all contingent propositions from the internal standpoint? Could we not so specify the internal standpoint as to admit more doxastic attitudes, even beliefs? Well, we certainly *could*, as many epistemologists have done by admitting "premisses" about the uniformity of nature, or the initial credibility or reliability of memory. But on what ground may we impute such doxastic attitudes to the internal standpoint? How are we to choose the particular doxastic attitudes? Presumably the selection should not be made randomly or arbitrarily. Should we choose beliefs that are widespread among human cognizers? This would be a poor rationale, since widespread beliefs may have no claim to epistemological priority. Should we impute beliefs that strike us as epistemically sound? That sort of strategy would seem to rely on *prior* epistemic standards, some prior commitment to a DDP. But the entire rationale for inventing the internal standpoint is precisely to *choose* a DDP (for the first time!). Basing such a choice on some antecedent DDP is indeed self-defeating and invites the prospect of a vicious regress.

Some readers may not be fully convinced by my attack on the feasibility of generating a unique DDP from the internal standpoint. They might feel that recent epistemologists in the foundationalist tradition have succeeded in producing doxastic principles of the sort I have in mind (if not an entire DDP), e.g., such epistemologists as Roderick Chisholm or John Pollock.[16] Are these epistemologists not good examples of internalists, and do they not succeed in displaying the favorable prospects for internalism?

Although Chisholm and Pollock may indeed be foundationalists, I do not think they are thoroughgoing internalists *in the sense of this notion that I have specified.* Whatever the merits of their work, they do not succeed in supporting our version of internalism.

Chisholm's epistemic principles may be viewed as partial DDPs. These rules assign epistemic status to propositions, and this assignment of epistemic status is roughly equivalent to prescribing, permitting, or prohibiting belief. To say in Chisholm's terminology that a proposition is "evident" for a person is to say, roughly, that

he ought to believe it. The crucial question is how Chisholm's principles are supposed to be derived. For example, Chisholm endorses the principle, "If you believe that you *perceive* something to have property F, where F is a sensible characteristic, then you ought to believe the proposition that there *is* something which has property F (before you)."[17] Chisholm would presumably be unwilling to endorse either (i) "If you believe that you *telepathize* a person to be thinking thought T, then you ought to believe the proposition that someone (else) *is* thinking thought T," or (ii) "If you believe that you *'clairvoy'* that x will happen, then you ought to believe the proposition that x *will* happen." Why is the principle concerning ostensible perception, but not the principles concerning ostensible telepathy or clairvoyance, to be endorsed? Chisholm proceeds by assuming that "we have at our disposal certain instances that the rules should countenance or permit and other instances that the rules should reject or forbid."[18] In other words, we rely on our commonsensical intuitions or judgments concerning what we know, or what the proper epistemic standards are.[19] This suggests that Chisholm's choice of epistemic principles rests on *prior* epistemic standards, which contravenes the idea of internalism. To put the point slightly differently, Chisholm seems to be engaged merely in codifying or systematizing antecedent doxastic practice, at least reflective doxastic practice. But it cannot be assumed that this practice, or the epistemic standards it expresses, derives from the internal standpoint. There is no guarantee, or even a hint of a guarantee, that our intuitive epistemic judgments could have been properly chosen from the internal standpoint. Hence nothing in Chisholm's discussion alleviates the difficulties for internalism.

Essentially the same point holds for Pollock. Pollock says that the "basic epistemological task" is to "*spell out*" the justification conditions for statements in different areas of knowledge.[20] This apparently means that epistemology should elucidate the set of justification conditions (roughly, the DDP) that we antecedently or commonsensically accept. True, Pollock maintains that the *meanings* of our ordinary physical-object statements, statements about the past, etc., consist in criteria of justification. But even if this is right, it just implies that we have chosen to use words in a way that commits us to certain justification principles (a DDP). There is no assurance, however, or even the glimmer of a suggestion, that such principles have been chosen, or could have been chosen, from the internal standpoint. So Pollock, like Chisholm, offers no comfort to internalism as *I* have delineated it.

VII

The epistemological difficulties of the Cartesian or egocentric position are familiar, and the previous section only reformulates these difficulties within a certain structure or framework. (I believe, however, that this framework makes the problems clearer than they have sometimes been made.) In particular, it is a familiar fact that it is difficult for the internal standpoint to license belief in *contingent* propositions. Still, the domain of logic might seem to be a different kettle of fish. After all, it is (allegedly) an *a priori* domain. Should the internal standpoint not be able to license doxastic attitudes toward propositions of logic? Should it not be able to generate *certain* doxastic principles (partial DDPs), viz., principles that derive from logic?

The relative tractability of these subjects is an illusion. First let us consider how to generate doxastic principles from logic. An initially attractive principle is this: "For any propositions Q and P, if Q logically implies P, and you believe Q, then believe P." There are two objections to this principle. First, as Gilbert Harman points out,[21] it may be best to abandon belief in Q under these circumstances rather than adopt a belief in P. Second, this principle violates our admissibility constraint on inputs (as we pointed out in section II). People cannot in general tell infallibly whether a putative logical implication really is one or not; hence this condition cannot serve as an antecedent of a doxastic principle.

The principle might be amended to read: "If you *believe* that Q logically implies P, and you believe Q, then believe P." This meets the admissibility constraint, but it does not meet Harman's objection. In fact, Harman's point is doubly troublesome here, since it is also problematic whether you should retain your belief in the putative logical implication; perhaps *that* belief should be abandoned.

A possible further amendment is this: "If you *justifiably* believe that Q logically implies P, and you *justifiably* believe Q, then believe P." This amendment might accommodate Harman's worry, but the resulting principle can serve only as a recursive principle and it focuses our attention on the need for *base clauses* concerning logical relationships (e.g., propositions of the form "Q logically implies P"). In light of our earlier restrictions (see section II), these base clauses must specify *non-epistemic* conditions in which one should believe logical truths.

Bayesians and confirmation theorists of various persuasions have often said that every tautology or logical truth should be assigned a subjective probability of 1.0. In our framework this amounts to the principle "For any proposition p, if p is a logical truth, then believe p (or be completely certain of p)." But this violates our "can tell" constraint on inputs. Ignoring that constraint for the moment, it should be clear that this principle is misguided. In terms of our regulative conception, it says that a person is automatically *justified* in believing (indeed, being certain of) any proposition that is a logical truth. Such a principle has no appeal. It conflates the modal or logical status of a proposition with its epistemic status. Bear in mind that the logical status of a proposition is a matter, like any other, in which mistakes and confusions are possible.[22] A person who is untrained or poorly trained in logic is even more likely to make mistakes, or to fail to "grasp" or "intuit" a logical relationship when one is present. Thus there is no plausibility in saying, of any person whatever, that whenever he entertains a proposition that is a logical truth, he is justified in believing it (indeed, in being completely certain of it).

Thus far the difficulties we have canvassed concerning matters of logic are as much a problem for externalism as for internalism. But our discussion highlights two peculiar difficulties for internalism. First, we have seen that propositions of logic are epistemologically problematic. Beliefs in such propositions, therefore, cannot be allowed in the internal standpoint; they cannot be used to choose a DDP. This exacerbates the difficulty of choosing a DDP from the internal standpoint, for it is hard to *figure out* what doxastic policies to adopt if one cannot employ (beliefs in) truths of logic in one's figuring. Second, our discussion reminds us of the fact that an appropriate doxastic strategy vis-à-vis propositions of logic partly depends on our powers

of reasoning and intuition. Well-chosen doxastic instructions should reflect the scope and accuracy of our imaginal and computational faculties. The exact nature of these powers and faculties, however, is a contingent matter, so there cannot be doxastic attitudes concerning them in the internal standpoint. This adds still another reason why no choice of doxastic principles for the domain of logic can be generated from the internal standpoint.

VIII

Internalists may object that I have not been entirely fair to them. Acknowledging the force of the difficulties I have raised, they may hasten to say that these difficulties stem largely from the assumption that true-belief-cum-error-avoidance is the proper goal of cognition. If a different goal is posited, it may be easier to choose a DDP from the internal standpoint.

Is there a plausible alternative goal? One classical candidate is *coherence* among doxastic attitudes. Coherence is a notoriously ambiguous concept, though. One version interprets it as *consistency*: a set of beliefs is coherent if and only if the propositions believed are logically consistent. A set of subjective probabilities is coherent if and only if they conform to the probability calculus. It has often been pointed out, however, that mere consistency (in either form) is an implausibly modest goal. It is too easy (relatively speaking) to obtain a merely consistent set of beliefs. Furthermore, it is unclear on this view why scientists and scholars regard the accumulation of additional evidence, and performance of new tests and experiments, as worthwhile intellectual enterprises. If mere consistency is our aim, new evidence and experiments are irrelevant. We can retain consistent beliefs with as sparse an evidence-base as you please. Evidently, we gather more evidence because more evidence (we assume) will contribute to greater truth-acquisition and/or error-avoidance.

Other forms of coherence have been advanced, *explanatory* coherence being the leading example. The notion of explanatory coherence, though, is far from clear.[23] Furthermore, it is hard to see why it should be regarded as the goal of cognition unless best explanatory hypotheses are likely to be *true*. So it is truth once again that surfaces as the fundamental cognitive desideratum.

Other candidate goals may be mentioned briefly. Peirce regarded the aim of inquiry as the "fixation" of belief, by which he meant the formation of belief in such a fashion that subsequent change is avoided. But why should mere fixity of opinion be sought? Besides, it is too easy a goal to meet, at least if doxastic voluntarism is true. One kind of DDP that meets this goal is a DDP that instructs cognizers to retain beliefs they have held in the past. Unfortunately, there are many different total DDPs that would contain this instruction. There is no way to choose from these DDPs, since they all meet the specified goal. This shows that the goal is too weak.

Another candidate desideratum out of the pragmatist tradition is the goal of "relieving agnosticism." Isaac Levi's treatment renders this equivalent to the goal of maximizing belief "content."[24] But Levi couples this with the goal of error-avoidance, and this again introduces an element that internalism finds difficult to handle. It is

clear, moreover, that content-maximization by itself is an implausible *exclusive* desideratum. A DDP could promote this goal by prescribing belief in every proposition, including contradictory propositions, which is radically counter-intuitive. Nor is it plausible to couple content-maximization with mere inconsistency-avoidance. There are indefinitely many consistent-but-detailed sets of possible beliefs. Which of these should be chosen? There would be no criterion of selection from these sets on the basis of this twofold goal.

I conclude that truth-acquisition and error-avoidance is the best candidate for the goal of cognition (or doxastic-attitude formation). Postulation of this goal does indeed pose problems for internalism, but that reflects ill on internalism, not on the goal.[25]

IX

We have surveyed numerous attempts to articulate the vague idea that underlies internalism. Most of these attempts proved unsatisfactory. When we finally formulated a condition that did not seem palpably inadequate, i.e., (6), we found good reason to believe that no unique DDP[26] could be generated from this condition. There seems, then, to be no way of fulfilling internalism's conception of the right DDP. Should we conclude that internalism is a mirage, a conception that beckons to us with no real prospect of satisfaction? Many epistemologists would shrink from this conclusion, since it would leave nothing but externalism to satisfy our need for a theory of justification. But I shall argue that we should be thoroughly content with externalism, that it offers everything one can reasonably expect (if not everything one has always wanted) in a theory of justification.

One objection to externalism is its failure to *guarantee* epistemic access to the optimal DDP, the DDP that externalism says is right. But this fact should not be confused with the claim that the optimal DDP is necessarily *in*accessible. Nothing in the conception of an optimal DDP precludes the possibility of successful identification of the optimal DDP. Furthermore, since two distinct total DDPs can share a large proportion of instructions, the uniquely optimal DDP may be approximated to a greater or lesser extent by numerous sub-optimal DDPs. We may get "close" to the optimal DDP, even if we do not get every detail exactly right. And if we regulate our doxastic attitudes in accordance with a slightly sub-optimal DDP, we may still "do" almost everything correctly, i.e., justifiably; for the great majority of our doxastic attitudes may conform with the right (= optimal) DDP as well as with our approximation to it.

Still, at any given time, we may not be in possession of the right DDP. Can this be reconciled with people's common practice of rendering epistemic appraisals? Certainly it can. People's appraisals of doxastic attitudes as justified or unjustified reflect only their *beliefs* about justifiedness, not the *facts* of justifiedness. When Smith judges Jones's belief to be justified (in the regulative sense), this indicates that Smith *believes* that Jones's belief accords with the right DDP. It does not mean that Jones's belief *does* accord with the right DDP, nor even that Smith is *justified* in believing that it so accords (though doubtless Smith *thinks* that he is justified).

However, many people have great confidence in their epistemic appraisals. Can this be reconciled with the externalist conception, which makes the right DDP difficult to identify? Part of the answer is that ordinary people have no inkling of the theoretical problems and difficulties and may therefore have *misplaced* confidence in their appraisals. With increased sophistication, some of this confidence may fade. The same is true of moral judgments. The mere fact that ordinary people often have great confidence in their moral judgments hardly shows that the correct moral standard (assuming there is one) is easily accessible. The common man's confidence in his moral appraisals may be misplaced.

But do not philosophers, logicians, and statisticians often have great confidence in their epistemic appraisals, which are not so readily dismissible as misplaced? Yes, but many of these appraisals are *negative*, and it is much easier to *fault* a piece of reasoning, or a conclusion, than it is to sustain one. If you can see that a certain pattern of thought will lead systematically to error, you can plausibly conclude that some other pattern of thought is better, and hence that the conclusion in question is not licensed by the optimal DDP. You can do all this without pretending to know exactly what the optimal DDP *is*.

Finally, even *positive* epistemic appraisals do not imply putative knowledge of the *total* optimal DDP. Saying that a certain doxastic attitude accords with the optimal DDP does not suggest that you can specify that DDP, any more than saying that a certain act conforms with (fails to violate) the criminal code implies that you can specify the details of that code in its entirety.

What is most perturbing about externalism, however, may be something entirely different. Suppose that the actually optimal DDP can be found only with difficulty and effort. *What are we supposed to do in the meantime?* How should we form our doxastic attitudes? How should we search for the optimal DDP? Which (meta-) DDP should we use in the search procedure?

The first point to note is that internalism faces the very same problems. Whatever the internalistically right DDP is supposed to be, we could not rely on its falling to us from heaven. We should probably have to work to get it. The same questions, then, arise: how should we form our doxastic attitudes in the meantime? and which DDP should we use in searching for the internalistically right DDP? These worries do not create a special presumption against externalism, since they are of equal significance for internalism.

Furthermore, there is an important reason why these questions and worries probably have no application, either to externalism or to internalism. This point rests on the probable falsity of doxastic voluntarism, at least the radical version of this doctrine which says that doxastic attitudes can *only* be chosen, that they cannot be formed by non-deliberative processes. On the contrary, there are native, or constitutional, doxastic processes that generate beliefs independently of our will and independently of the deliberate selection of a DDP. Perceptual processes automatically produce representations that, unless inhibited by other cognitions, serve as beliefs. Similarly, we are all ground-level inductivists. Expectation based on past experience is part of our animal heritage. Thus we do have means of forming doxastic attitudes

before choosing doxastic *principles*. Native doxastic *habits* render the selection of doxastic *principles* not strictly necessary. When we do come to choose doxastic principles, this choice may be based on beliefs formed by antecedent doxastic habits. No meta-DDP is required.

Indeed, a stronger argument can be made. Not only is it *possible* for the choice of a doxastic principle to be based on doxastic habit; it is *necessary*. If every choice of a DDP rests on doxastic attitudes that result from a prior DDP choice, there would be an infinite regress of DDP choices. Such an infinite regress is impossible (at least physically impossible). So if there are any DDP choices at all, they must rest, ultimately, on doxastic habits.

This conclusion jibes with the most plausible reconstruction of how doxastic principle choices come to be made. At the start a creature forms beliefs from automatic, preprogrammed doxastic processes; these beliefs are largely about its own immediate environment. At a later stage of development, at least in a sophisticated creature, beliefs are formed about its own belief-forming processes. The creature comes to believe that certain of its belief-forming processes often lead to error and that others are more reliable. This kind of belief can arise as follows. The creature predicts a certain event, i.e., believes it will occur. The creature then observes its non-occurrence, i.e., believes it has failed to occur. If the creature remembers the earlier belief, it can conclude that the earlier belief was false, hence that the process which generated it is (somewhat) unreliable. These beliefs presuppose, of course, that the creature has a conception of truth and falsity. This need not be a philosophical *theory* of truth but merely a rudimentary conception thereof. (Such a conception would seem to go hand-in-hand with a conception of belief, since believing a proposition is believing it to be true.) Once the creature distinguishes between more and less reliable belief-forming processes, it has taken the first step toward doxastic appraisal. It can then introduce a (non-regulative) notion of *justifiedness*; beliefs are justified if and only if they are produced by (relatively) reliable belief-forming processes. The creature can also begin doxastic self-criticism, in which it proposes *regulative* principles to itself.[27] On the assumption that doxastic states are at least partly subject to deliberate decision, the creature can formulate and endorse DDPs that might avert errors or biases in its native processes.

What is noteworthy in this reconstruction is that the notion of justifiedness, especially the regulative notion, is a late arrival. It does not appear until the creature already has beliefs and a conception of truth. This reverses the scenario that internalism tends to foster. Internalism encourages the idea that the choice of a DDP should antecede all belief, that we must first select a *criterion* of truth—a principle for deciding which propositions are true—before we form any beliefs. But if my description of how things actually go is correct, this order is reversed. Beliefs come before the selection of a DDP. Upon reflection, it is hard to see how things *could* go otherwise, given the kind of creature we are.

If the foregoing discussion is right, the answers to our earlier questions are straightforward. We do have a means of forming doxastic attitudes prior to our (initial) choice of a DDP: we have our native, preprogrammed doxastic processes.

Similarly, there is a straightforward answer to what we should do in trying to identify the optimal DDP. At first we should (and must) use our constitutional doxastic habits. Once these habits generate the choice of a DDP, that DDP should be used (together with the habits which it does not wholly displace) to form any new views about the optimal DDP.

Actually it is an open question whether *any* doxastic processes can be deliberately influenced. If all doxastic processes are, in Hume's terminology, "permanent and irresistible," then the entire conception of a DDP has no application whatever. A middle ground is possible, though. *Some* doxastic processes may operate independently of the will and may not be modifiable by voluntary means: *other* doxastic processes may be subject to reflective and voluntary direction. In this case—which is the most convenient one for present purposes—there would be points of application for DDPs, but there would also be automatic processes that operate prior to the appearance of reflective choice in the doxastic domain.

Notice that even if the *doxastic* domain does not admit of voluntary intervention, there is still scope for deliberate self-guidance in the intellectual sphere. There are doubtless other events or states—distinct from doxastic states themselves—that can be voluntarily controlled and have indirect influence on doxastic attitudes. The search for evidence (e.g., deployment of the sensory apparatus in purposefully selected directions) and the search for, or attention to, hypotheses on deliberately chosen topics are examples of voluntary processes that have a major impact on our beliefs. So even the falsity of *doxastic* voluntarism would not undermine the regulative orientation of epistemology.[28] However, we might have to substitute the idea of a CDP—a *cognitive* decision principle—for that of a DDP—a *doxastic* decision principle.

The perspective advanced in our recent discussion has salient points of contact with epistemological "*contextualism*," e.g., the version of this approach adopted by W. V. Quine. Quine emphasizes that there is no point of cosmic exile.[29] We have to start, epistemologically speaking, from the beliefs we have at a given time. Similarly, Karl Popper emphasizes that rational inquiry can consist only in criticism of antecedently held belief. We do and must start from an uncritically held body of opinion, or mental structure.[30] The same view is expressed by C. S. Peirce: ". . . there is but one state of mind from which you can 'set out,' namely, the very state of mind in which you actually find yourself at the time you do 'set out'—a state in which you are laden with an immense mass of cognition already formed, of which you cannot divest yourself if you would . . .".[31]

My contextualism, however, is only a view at the *meta-theoretical* level. It only answers the question: How should we go about choosing a criterion of justifiedness, i.e., a DDP? My answer to *this* question is entirely neutral with respect to the specific criterion of justifiedness that should ultimately be chosen, i.e., the optimal DDP. That DDP could turn out to be some sort of "*foundationalist*" DDP. Its belief-prescriptions might reflect the idea that there is a special class of propositions that should be believed quite independently of other beliefs and from which all other beliefs should be inferred. I do not mean to *endorse* this sort of foundationalism.

I cite it only to illustrate the point that, whereas Quine, for example, is a thoroughgoing opponent of foundationalism, my own *two-leveled* theory allows more room for epistemological flexibility. On the meta-level—"How can we try to identify the right DDP?"—I take a contextualist position. But on the basic level—"Which DDP is the right DDP?"—any answer is viable (in principle).[32]

The reader may feel that this very flexibility is a weakness of my discussion. If I have not specified the right DDP, or even the general form of the right DDP (foundationalist, coherentist, etc.), what has been accomplished? Two important things, I believe, have been accomplished. First, we have elucidated the general idea of a regulative conception of justifiedness and shown how justificational status, so conceived, depends upon (a) the cognizer's current states and (b) the right DDP. Second, we have articulated the only tenable conception of DDP-rightness (externalism) and disposed of an illusory rival conception (internalism) that traditional epistemology suggests. We have thereby provided two fundamental building blocks of any fully detailed regulative epistemology, at least any regulative epistemology that seeks to guide *doxastic* choices.[33]

Notes

1. In A. I. Goldman and J. Kim, eds., *Values and Morals* (Dordrecht, 1978). See p. 194.

2. "Moral Decision Principles."

3. "What is Justified Belief?," in George S. Pappas, ed., *Justification and Knowledge* (Dordrecht, 1979).

4. One general problem with this approach is the assumption of a uniquely correct DDP. Suppose instead that two or more DDPs are "*tied*" for best and that some of these make conflicting prescriptions. In particular, suppose one DDP tells S to believe proposition p and another tells S to disbelieve it. Which doxastic attitude vis-à-vis p is S justified in adopting? The solution is to say that a cognizer is justified in having doxastic attitude D vis-à-vis p at *t* if and only if there is *at least one best* DDP that prescribes D vis-à-vis p at *t*. This is tantamount to saying that a cognizer is justified in adopting D vis-à-vis p just in case he is *permitted* to have D vis-à-vis p and that he is permitted to have D vis-à-vis p if at least one best DDP prescribes it. This is indeed quite plausible. In fact, apart from the problem of multiply correct DDPs, there is much to be said for linking justifiedness with a *permission* to adopt the indicated doxastic attitude rather than a *prescription* to do so. Even a single DDP may contain permissions as well as prescriptions and prohibitions, and it is plausible to say that a doxastic attitude is justified if it is permitted.

For purposes of simplicity, however, I am going to bypass these issues. First, I shall link justifiedness to what is prescribed rather than to what is permitted. Second, I shall generally assume that there is a uniquely right DDP. More precisely, I shall ignore the problem of "ties." In later discussion when I *do* admit the possibility of there not being a uniquely correct DDP, I shall have in mind not the possibility of ties but the possibility of correctness being relativized to different cognizers.

5. Is this explication of "can tell" too strong? Admittedly, it may be controversial whether even current mental states meet the requirement as formulated. In all likelihood, such states as *stored* beliefs, especially the *totality* of one's stored beliefs, do not satisfy this constraint. But we can live with such consequences. The regulative conception need not try to decide precisely which states (if any) satisfy its constraint on inputs. This would only have to be decided by anyone seeking to *use*, or *apply*, the regulative conception. For simplicity I am just assuming that a variety of current cognitive states do satisfy the constraint.

Still, we can entertain the prospect of weakening the explication in the text. But something like it seems needed. A principle cannot be a genuine decision-principle unless a person can

actually be guided by it, i.e., act in conformity with it. He must be able to bring it about that he executes an output-prescription when and only when the corresponding input-condition is fulfilled. To do this, it seems, he must have the power to tell, or detect, when each input condition (antecedent of a principle) is or is not fulfilled.

6. More precisely, we are not interested in a recursive "account" of justification but in a recursive DDP, i.e., a recursive set of doxastic principles. The consequent of a doxastic principle will not be something of the form "you are justified in believing Q," but rather something like "believe Q." Hence, perhaps a better schema for an inductive (or recursive) member of a recursive set of principles would be this: "If, given your present inputs, this DDP instructs you to believe Q, and Q bears such-and-such a relation to P, then believe P."

7. The indicated goal may bias the choice of a DDP in favor of "belief"-principles as opposed to "subjective-probability"-principles. To avoid this bias we could elaborate the goal by giving positive value to assignments of subjective probabilities greater than .5 to truths and negative value to assignments of subjective probabilities greater than .5 to falsehoods. In the remainder of our discussion, however, I shall ignore this complication.

8. It is part of the regulative conception in general—hence common to both externalism and internalism—that a cognizer must be able to tell which *input* conditions obtain. But here we are discussing *DDP-rightness*, not input conditions. The regulative conception *per se* imposes no requirement of epistemic access to the right DDP. It merely says that a doxastic attitude is justified if it is prescribed by the DDP that is *in fact* right, whether or not the cognizer knows it is right. Internalism goes further by seeking to make it a condition of a DDP *being right* that a cognizer can know or tell that it is right (or optimal).

9. If relativization is permitted, we need a slight amendment in the basic relationship between justifiedness and DDP-rightness. That relationship would then be formulated as follows: S is justified in having doxastic attitude D vis-à-vis p at *t* if and only if the right DDP (*or the DDP that is right for S at t*), when applied to the relevant input conditions that characterize S at *t*, yields as output the prescription ⌜adopt attitude D vis-à-vis p⌝.

10. For a distinction between objective and subjective justification in epistemology, see John L. Pollock, "A Plethora of Epistemological Theories," in Pappas, ed., *Justification and Knowledge*. It is arguable that Keith Lehrer's view in *Knowledge* (London, 1974) is a form of subjectivism or relativism. On the other hand, it might be construed as a kind of non-relativistic coherentism.

11. A more difficult case is proposed by Holly Goldman, Suppose that a cognizer uses the right DDP—call it R—to decide which DDP is right, and R prescribes belief in the *false* proposition that DDP W is right. Then this cognizer is *justified* in believing that W is right. (Suppose also that W prescribes belief in its being right.) Now suppose he considers whether to believe proposition p, and W prescribes belief in p. On the other hand, R prohibits belief in p. Is the cognizer justified in believing p? According to the basic structure of the regulative approach (see section I), he is not so justified, since the right DDP does not prescribe belief in p. On the other hand, there is a strong temptation to say that he *is* justified in believing p. After all, he is justified in believing that he is justified in believing p, and it is an attractive principle that if one is justified in believing that one is justified in believing p (JJp), then one is justified in believing p (Jp). However, we are not forced to accept this principle, and an epistemological objectivist is well advised not to accept it.

12. However, it is not clear that all traditional epistemologists have tried to *select* a DDP from the internal standpoint. Many have simply presupposed some (inadequately specified) DDP and explored the question whether this DDP, when applied to the internal standpoint, prescribes our commonsensical beliefs, e.g., belief in the external world. Thus, what I am here calling "internalism" may not correctly describe most historical epistemology. Nonetheless, if epistemology accepts the doxastic regulative enterprise, the selection of a DDP is absolutely central, and criteria for the rightness of a DDP are equally crucial. Once this problem is posed, it seems consonant with the epistemological tradition to suggest that DDP-rightness be judged from the internal standpoint.

13. Indeed, there is not even a *small group* of DDPs that can be picked out as best.

14. It might be thought that the regulative conception itself automatically confers special status on first-person-current-mental-state propositions, for the regulative conception (RC) presupposes that cognizers can infallibly form beliefs about their own current mental states. We should avoid misunderstanding, however. The RC just lays down a *constraint* on input conditions (the antecedents of DDPs), i.e., that they must meet the "can tell" requirement. The RC does not say which *specific* conditions meet this requirement, if any. *I* claimed that, plausibly, a person's cognitive states—at least many such states—meet this requirement. But the RC per se makes no commitment on this point. In any case, even if it is *true* that such-and-such mental states meet this requirement, it does not follow that beliefs in propositions describing these states should be allowed in the internal standpoint. These propositions may still be epistemologically problematic.

15. It is doubtful that maximax would select a unique DDP. There are infinitely many incompatible DDPs that prescribe equally profligate belief-formation. The best outcomes of each of these will be equally good.

16. See Roderick M. Chisholm, *Theory of Knowledge* (Englewood Cliffs, N.J., first edition, 1966, second edition, 1977) and John L. Pollock, *Knowledge and Justification* (Princeton, 1974).

17. The formulation here includes some rewording of my own.

18. Chisholm, *Theory of Knowledge*, first edition, p. 24.

19. See *ibid.*, second edition, p. 16.

20. *Op. cit.*, p. 21.

21. Gilbert Harman, *Thought* (Princeton, 1973), p. 157.

22. This is stressed by Hume, among others. See *A Treatise of Human Nature*, book I, part IV, section I.

23. See Lehrer, *Knowledge*, chap. 7.

24. Isaac Levi, *Gambling with Truth* (New York, 1967).

25. Still another possible goal is the satisfaction of needs or desires. But it is not easier for internalism to generate a DDP for this goal than for the goal of true-belief-cum-error-avoidance. Indeed, need or desire satisfaction will depend heavily upon truth-acquisition and error-avoidance.

26. Or even any small set of best DDPs.

27. I do not mean to imply here that the non-regulative conception of justifiedness must temporally precede the regulative conception.

28. This point is also made in "Epistemics: The Regulative Theory of Cognition," *The Journal of Philosophy*, 75, no. 10 (October 1978):509–23.

29. W. V. Quine, *Word and Object* (Cambridge, Mass., 1960), p. 275.

30. Karl Popper, *Objective Knowledge* (Oxford, 1972), pp. 71–72.

31. Justus Buchler, ed., *Philosophical Writings of Peirce* (New York, 1955), p. 256.

32. One might think that the regulative conception (RC) is biased in favor of foundationalism. Is it not committed to the idea that we have infallible access to our own cognitive states? And is this idea not peculiar to foundationalism? Two comments are in order. First, as indicated in notes 5 and 14 above, the RC does not say *which* states of a person, if any, have this special status. It merely lays down a *constraint* on inputs. Second, it is not true that infallible access to cognitive states is a thesis peculiar to foundationalism. Coherentism is committed to the same sort of thesis. In particular, it is committed to the assumption that we can tell what our own current *beliefs* are. Unless we can tell what (all) our current beliefs are, we cannot tell whether they satisfy the requirement of coherence (whatever exactly this requirement is), nor can we adjust our beliefs to meet that requirement.

33. Some of the ideas in this paper germinated during research that was sponsored by the John Simon Guggenheim Foundation and the Center for Advanced Study in the Behavioral Sciences. For many helpful suggestions and criticisms I am indebted to Holly Goldman, Louis Loeb, and Terence Horgan.